




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ :	A1	(11) International Publication Number:	WO 98/28938
H04Q 11/04, H04L 12/56		(43) International Publication Date:	2 July 1998 (02.07.98)

(21) International Application Number:	PCT/CA97/00936	(81) Designated States: CA, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).
(22) International Filing Date:	3 December 1997 (03.12.97)	Published <i>With international search report.</i>
(30) Priority Data:	23 December 1996 (23.12.96) US 08/772,256 14 March 1997 (14.03.97) US 08/818,612	
(71) Applicant:	NORTHERN TELECOM LIMITED [CA/CA]; World Trade Center of Montreal, 8th floor, 380 St. Antoine Street West, Montreal, Quebec H2Y 3Y4 (CA).	
(72) Inventors:	CHAPMAN, Alan, Stanely, John; 50 Pentland Crescent, Kanata, Ontario K2K 1V5 (CA). KUNG, Hsiang-Tsung; 1 Joseph Comee Road, Lexington, MA 02173 (US).	
(74) Agent:	TOYOOKA, Yoshiharu; Northern Telecom Limited, Patent Dept., P.O. Box 3511, Station "C", Ottawa, Ontario K1Y 4H7 (CA).	

(54) Title: DYNAMIC TRAFFIC CONDITIONING



(57) Abstract

Multi-media networks will require that a data flow be given certain quality-of-service (QOS) for a network connection but pre-negotiation of this sort is foreign to the current data networking model. The real time traffic flow in the data network requires distinct limits on the tolerance to delay, and the variations in that delay. Interactive voice and video demand that the total delay does not exceed the threshold beyond which human interaction is unacceptably impaired. The present invention allows the network to discover the nature of the service for each traffic flow, classifies it dynamically, and exercises traffic conditioning by means of such techniques as admission control and scheduling when delivering the traffic downstream to support the service appropriately.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		

DYNAMIC TRAFFIC CONDITIONING

This is a continuation-in-part of Application Serial No. 08/772,256, filed on December 23, 1996.

5

Field of the Invention

The invention relates generally to the conditioning of traffic at a node of a data network. In particular, it is directed to a technique by which the traffic is dynamically classified at a node and delivered 10 downstream according to the Quality of Service (QOS) specified by the assigned class or by such a network administrator as the network operator.

Background of the Invention

15 Multi-media networks will require that a data flow be given certain QOS for a network connection. The recently proposed resource reservation protocol (RSVP) for IP (and signalling in ATM networks) is a way of requesting a particular QOS but pre-negotiation of this sort is foreign to the current data networking model and would require 20 changes at the application level.

The major driving force behind the requirement for different QOSs in the data network is the need to introduce real time flows which have distinct limits in the tolerance to delay, and the variations in that delay. Interactive voice and video demand that the total delay 25 does not exceed the threshold beyond which the human interaction is unacceptably impaired. Non-interactive voice and video streams which are being transferred in real time require that the maximum delay variation is bounded so that buffers can be kept to a reasonable size and guaranteed not to underflow. Meeting the delay requirements for real time flows usually means that these flows must be given 30 priority over other traffic. This in turn brings in a requirement to limit the amount of such high priority traffic by some admission control policy to ensure that other classes of traffic do get some of the available bandwidth.

35 There are applications other than voice and video that can benefit from controlled latency. Network control traffic such as DNS

transactions represent a small fraction of the total but will provide a much improved performance if treated with priority.

There is another class of traffic which does not have the tight requirements of voice or video but does involve human interaction

- 5 with computers and can lead to noticeable decreases in productivity (or increases in frustration) if subjected to long delays. This traffic type is generated by applications such as X-Windows, Telnet and, more often now, world wide web browsing. This traffic can be protected from long queuing delays caused by bulk transfers such as FTP or NFS by
10 allocating to it some guaranteed portion of the bandwidth as part of an output scheduling policy.

Even the bulk traffic can suffer from too much competition.

Often a file transfer will be aborted after using much network resource because the overall time has exceeded the delay tolerance of the

- 15 application or the user or management policies in intermediate servers. Also, when congestion causes packets to be dropped, it can easily impact many flows, and cause many resends. By guaranteeing a certain number of flows a minimum bandwidth and treating the remainder as best effort, it is possible to avoid spreading packet loss
20 over so many flows and to reduce the number of aborted flows.

It would be much more acceptable if the QOS requirements were met by the network automatically and dynamically without the need for signalling. This would fit more naturally with the current IP networking paradigm.

- 25 Traditionally, Internet services (such as FTP, Telnet, NFS) are known only to the end systems and not to the network itself. The present invention allows the network to discover the nature of the service for each traffic flow, classify it dynamically and exercise traffic conditioning by means of such techniques as admission control and
30 scheduling when delivering the traffic downstream to support the service appropriately. The scheduling separates real time traffic from other traffic by priority and allocates bandwidth between various classes of traffic. In conjunction with scheduling, the admission control guarantees performance. The scheduling also allows implementation
35 of overlay administrative policies to give, for instance, certain groups different treatment than other groups. The classification need not emulate precisely the effect of pre-negotiated network connections but

should provide similar improvements in service quality as seen by the users and the network.

It should of course be noted that in this specification the data network can also include any packet-based or cell-based networks,
5 including ATM networks.

Objects of Invention

It is therefore an object of the invention to provide a method of and an apparatus for dynamically conditioning traffic at a node of a
10 data network.

It is a further object of the invention to provide a method of and apparatus for continuously monitoring and dynamically classifying traffic into one of a plurality of preset classes according to a set of classification parameters.

15 It is another object of the invention to provide a method of and apparatus for controlling the delivery of the traffic downstream according to quality of service parameters specified by the dynamically selected class.

20 Summary of the Invention

Briefly stated, the invention resides in a packet data network for multimedia traffic having one or more nodes. According to one aspect, a method of dynamically conditioning traffic comprises steps of continuously monitoring the traffic at a node as to its traffic

25 characteristics and classifying the traffic into one of a plurality of classes according to preset criteria of traffic characteristics. The method further includes a step of delivering the traffic downstream according to the quality of service specified by one of the plurality of classes.

According to another aspect, the invention is directed to a
30 method of dynamically conditioning traffic comprising steps of continuously monitoring the traffic at a node to detect a real time traffic flow and determining if the real time traffic flow can be admitted for delivery downstream, based on the available bandwidth of the data network. The method includes another step of reclassifying the real time traffic into a different class for delivery downstream according to a quality of service specified by the different class.

Brief Description of Drawings

Figure 1 is a schematic diagram of the traffic conditioner according to an embodiment of the invention.

Figures 2 and 3 show possible locations of traffic conditioners of
5 the invention.

Figure 4 illustrates that the traffic conditioning features of the invention can be installed at a router or a switch.

Figure 5 is a pictorial view of state transitions.

10 Figure 6 shows classification information being propagated downstream.

Figure 7 shows that the REJECT information is being passed forward and reflected back to the source host.

Figure 8 shows a circumstance in which a stop message is sent back upstream from a node.

15

Detailed Description of the Preferred Embodiments of the Invention

Referring to Figure 1, the traffic conditioner, according to one embodiment of the invention, includes a plurality of queues 10, at least one for each class. It is located at a node of a data network. Every

20 packet of an input stream is inspected and identified at 12 using, for example, IP addresses, ports and protocol. A controller 14 characterises the flow (using rate, duration, etc.) and assigns it a class. A plurality of classes are envisioned according to embodiments of the invention and will be described later. The controller refers to a database 16 and uses
25 output scheduling to allocate bandwidth among classes and implements an admission control policy for a class before delivering an output stream toward downstream nodes or to peripherals. It also discards packets from non-admitted flows and passes on flow classification to downstream nodes.

30 The traffic conditioning of the invention can take place at various places in the data network. For example, as gateways are often a bottleneck and bulk flows can decrease response times for interactive users, a traffic conditioner can be located at a place shown in Figure 2 which will alleviate this problem. In Figure 3, traffic conditioners are
35 located at a plurality of IP switches which form a data network 40.

Packet switches will not be able to provide good performance for new services, such as real time video, unless traffic conditioning is

implemented. A conditioner at an output port of a switch 42 monitors and controls all the traffic which comes into the switch and goes out of this output port to the downstream node 44. An integrated traffic conditioning feature can be installed at a router or a server as shown in
5 Figure 4.

In one instance of IP networking, the identifying individual traffic flow is simple. It only requires inspection of addresses and protocol port number for both source and destination. This approach has been used for some time in other applications such as TCP Routing
10 and, more recently, in IP switching.

Any scheme which characterises and classifies flows for differing treatment should be capable of modification to suit the particular operating environment. According to one embodiment, the following six classes of traffic flow which could be used in a typical TCP/IP based
15 network are considered and described in detail below.

(Class 1) – Interactive Users (TCP)

This class is intended to capture TCP flows for such applications as X-windows, Telnet and lightweight web browsing. Telnet and X-windows can be very long sessions but have predominantly short packets. In this embodiment, this classification defines a short packet as less than or equal to 128 Bytes and a long packet as greater than 128 Bytes. Web browsing has a mixture of packet sizes and transaction lengths and can cause flows of long packets, particularly while loading
20 image files. This class is the default one for TCP flows but when the flow contains a long series (e.g., 200) of long packets without an intervening series of two or more short packets, it is considered to be bulk transfer and no longer valid for this class. A TCP bulk flow may occasionally include short packets, containing only the TCP ACK
25 message or the tail end of a sending buffer. Thus, unless the short packets are in a series of two or more in number, the flow is still considered to be bulk transfer.
30

(Class 2) – Bulk Transfer with Guaranteed Bandwidth (TCP)

If TCP traffic is not interactive, it is classified as bulk transfer. Some of the bulk transfer flows are admitted to a scheduling class
35 which has a protected portion of the bandwidth and a limited number

of members such that some minimum bandwidth is guaranteed for these flows.

(Class 3) – Bulk Transfer, Best Effort (TCP)

- 5 Any bulk TCP flows which are not admissible to the guaranteed bandwidth class are treated in this class on a best effort basis. Classes (2) and (3) are intended to handle large FTP or large web images.

(Class 4) – Low Latency (UDP)

- 10 This class contains flows which require very low bandwidth. In general, these flows consist of low bandwidth voice, NFS requests, short NFS replies, and network control packets such as DNS transactions. Most real time voice will be continuous flow UDP (User Datagram Protocol) at less than 20 packets per second. This is the
15 default class for UDP and flows above a threshold bandwidth are moved out fast. There would be an admission control policy such that if the allocated bandwidth for this class is used up, new flows would be moved to the best effort class.

20 (Class 5) – Real Time (UDP)

- Any flow with real time characteristics which will be defined below and a bandwidth too high for the low latency class will be captured in this class. High bandwidth voice and streaming video are the expected members. For example, video will be a continuous UDP
25 flow at up to 4Mb/s = (1000pps). Flows exhibiting real time characteristics but for which there is no resource under the admission control policy, would be rejected and subsequent packets discarded.

(Class 6) – Bulk Best Effort (UDP)

- 30 All other UDP flows will be grouped into this class. Expected members are NFS file copy and backup sessions characterised by many long packets at a substantial rate.

- In summary, according to this embodiment, discrimination between traffic types is based on simple analysis of packet arrival rates
35 and packet length plus a particular test for real time UDP traffic which will be described below. The characteristics of a flow are monitored continuously and the flow can be re-classified during its lifetime. The

ability to re-classify also covers the cases where consecutive flows assume the same identity but have different characteristics.

It is possible to use port number information in some instances to determine service type and complement the dynamic classification 5 described above. For example, TCP port 23 is the well known port for the telnet server. Since a telnet session always carries interactive traffic, it is possible to classify any TCP flow with one of its port number being 23 as TCP interactive. However, a single port number can sometimes be used for multiple services of varying nature. For 10 instance, the TCP port number 80, reserved for the HTTP server, could be used for both TCP Interactive and TCP Bulk. A classification scheme should use static information in the packet header such as port number and protocol types, as well as the information collected dynamically.

Figure 5 is a pictorial view of state transitions involving these 15 classes according to the embodiment. In the figure, the following abbreviations are used:

P = packet
sp = short packet
lp = long packet
20 con lp = consecutive long packets
bw = bandwidth available
chosen = chosen by policy
pps = packets per second arrival

The following criteria are also applied:

- 25 (a) For all classes, return to initial if there are no packets for 30 seconds.
(b) For all flows, check option field and if class is defined then force the flow to that class for 100 packets or 30 seconds.

Referring to Figure 5, the state machine contains the following 30 states:

- (A) Initial
- (B) TCP Interactive
- (C) TCP Bulk Best Effort
- (D) TCP Bulk Guaranteed
- 35 (E) UDP Low Latency
- (F) UDP Real Time
- (G) UDP Bulk Best Effort

(H) Reject

In the following detailed description, ">" indicates "true" for all the states.

5 Power up

> [Idle for 30 seconds]--> Initial

If class information is received from upstream then force that state for 100 packets. After 100 packets or 30 seconds with no information,

10 revert to local judgement.

> Initial (at 50 in Figure 5)

> [Protocol=TCP] --> TCP Interactive

> [Protocol=UDP] --> UDP Low Latency

15

> TCP Interactive (at 52)

> [200 consecutive long packets each of size > 128 bytes without an intervening series of two or more short packets]

> [Some TCP Bulk min bandwidth available]

20

--> TCP Bulk Guaranteed

> [Otherwise]

--> TCP Bulk Best Effort

> TCP Bulk Best Effort (at 54)

25 > [Departure of a flow from TCP Bulk Guaranteed] and [Picked by a Policy]

--> TCP Bulk Guaranteed

> [Two consecutive short packets of size no greater than 128 bytes]

30

--> TCP Interactive

> TCP Bulk Guaranteed (at 56)

> [Two consecutive short packets of size less than 128 bytes]

--> TCP Interactive

35

> UDP Low Latency (at 58)

> [After any 1 second in the state, and > 25 packets have arrived]

- > UDP Bulk Best Effort
- or
 - > [Two or more consecutive long packets] and [>5 packets in 1 second]
- 5 --> UDP Bulk Best Effort
 - >[Bandwidth is known] and [no more bandwidth available]
 - > UDP Bulk Best Effort
- > UDP Bulk Best Effort (at 60)
- 10 > [Matches real time template] and [bandwidth available in UDP video]
 - > UDP Real Time
 - > [Otherwise] --> Reject
 - > [Idle for more than one second] and [bandwidth available]
- 15 --> UDP Low Latency
- > UDP Real Time (at 62)
 - > [300th packet arrives in time < 200 ms]
 - > UDP Bulk Best Effort
- 20 > [Idle for 1 second]
 - > UDP Low Latency
- > Reject (at 64)
 - Discard all packets except those used for propagating REJECT information.

Real Time Flow Detection

As described earlier, it is necessary to distinguish real-time traffic from other BULK UDP flows. It is reasonable to assume at the present state of the art that NFS represents the only UDP traffic which is bulk and not real-time. Then an easy way is to inspect the port numbers of the flow in question. If neither of them is 2049, the reserved port number for NFS servers, then the flow is classified as Real Time.

Similarly, if it can be assumed that real-time flows of interest all use servers with reserved port numbers, then they can be classified by the static information on port numbers. Today, most of real-time

audio and video servers appear to use a set of special ports which can be easily recognized.

For circumstances where these assumptions don't apply, it is necessary to provide a dynamic method to identify real time flows.

- 5 The major difference between real time UDP flows and other UDP applications such as NFS is that (a) real time streams are not self-clocked, that is, the packets are sent continuously without acknowledgements from the receiver, and (b) the average packet generation rate is a constant.
- 10 There are two ways in which these attributes can be used to detect the real time flow. Under congestion conditions, the queue for a real time flow will grow without bound whereas the queue for a self-clocked flow will only grow to the size of an acknowledged burst (the maximum burst without acknowledgement is 8 Kbytes for NFS).
- 15 Another approach is to keep a history of interarrival times for the packets of a flow. Self-clocked flows will exhibit a bi-modal distribution whereas the real time streams will be uni-modal around the average rate. In one embodiment, this approach can be implemented by using a threshold that is set to a value somewhat more than the average
- 20 interarrival time. Two counts are kept; one of interarrival times below the threshold and one for interarrival times above the threshold. If the second count is more than a certain value e.g., 10% of the first count then the flow is classed as non-real time. This second method may be preferred since it does not require waiting for congestion to identify the
- 25 stream. Quick admission or rejection of the flow is important.

Scheduling Flows

There are various techniques available for scheduling multiple queues at output ports of a switching node. In any case, typically, the real time and low latency classes described above would be given absolute priority (hence the need for admission control), and the other classes would be scheduled in a way which allocates to them the proportion of the bandwidth allocated by the network administrator.

- 35 Referring to Figure 6, in a further embodiment, instead of each node performing the above classification autonomously as described thus far, the first node (node A) on the route passes on its findings to the downstream nodes and provides more efficient traffic

- conditioning. In yet a further embodiment, there is no reason that the source host machine cannot be considered the first node which classifies the traffic and sends the classification to nodes along the route. These downstream nodes can then use the acquired knowledge
- 5 instead of local classification and give the flow a consistent treatment throughout its path. They accept the upstream classification until it expires and propagate it downstream. The information can be propagated in various ways, but in one way it could be carried in the flow by inserting information in one or more packets in the flow (for
- 10 example, by an entry in the IP option field). Every n-th packet could carry the information and the downstream node would act on that information until it aged out after a multiple of n packets. In an ATM based network, the classification could be implied by the VPI/VCI value chosen for the flow.
- 15 Referring to Figure 7, the classification REJECT can be reflected back from the destination host to the source host, thus improving the efficiency of the network. To facilitate this, not all packets from the flow would be discarded. Instead, every k-th packet would be passed onward with the classification in the option field. For example, node B
- 20 determines that the packet must be rejected for whatever the reason, it discards all but every k-th packet, which is forwarded to the destination host, which in turn sends a message to the source host to stop. In another scenario as shown in Figure 8, instead of the host, a node which rejects the traffic stream for some reason can send a message
- 25 back upstream indicating that the traffic has been rejected and should therefore be stopped.

Complexity

- In any network where different qualities of service are being
- 30 supported, there has to be some process to classify the flow attributes and then a way of verifying that the flow is adhering to the attributes. This verification, or policing, can be done at the edges of a subnet or at every node. The method described here uses the mechanisms that would be in place for verification to discover the flow attributes but
- 35 without the necessity for pre-negotiation such as signalling. It has, therefore, an equivalent complexity in implementation at the nodes but does not require the signalling overlay.

The administrative overlay required for policies on bandwidth allocation to the various classes, and reservation of bandwidth for pre-booked video sessions, would be no different in complexity to a signalling based system.

5

Co-existence With RSVP and Administered Connections

Although the dynamic classification can run without the need for signalling, it is also possible to force the classification of the flow as a result of higher level processes. A flow that has been given a reserved path and treatment through the network by signalling or by administration can be marked as such. The automatic classification can be disabled or used to verify the characteristics of the marked flow.

According to the present invention, packet traffic flows are classified in order to group them for differing treatments. This allows quality of service distinctions to be supported even when application signalling support is not available. The implementation is no more complex than that required for policing when signalling is used and may be just as effective in improving perceived network performance and enabling new services such as video.

WHAT IS CLAIMED IS:

1. In a packet data network for multimedia traffic having one or more nodes, a method of dynamically conditioning traffic comprising
5 steps of:
 - continuously monitoring the traffic at a node as to its traffic characteristics;
 - classifying the traffic into one of a plurality of classes which includes two default classes, TCP interactive and UDP low latency,
 - 10 according to preset criteria of traffic characteristics which includes TCP protocol type, UDP protocol type, the length of packets in the traffic and the number of consecutive packets in the traffic, said plurality of classes specifying respective levels of the quality of service with which to condition the traffic; and
 - 15 changing the TCP interactive class to one or the other of TCP bulk guaranteed and TCP bulk best effort, if the traffic contains a series of long packets without an intervening series of two or more short packets.
-
- 20 2. The method according to claim 1 comprising a further step of:
 - changing the UDP low latency class to one or the other of UDP real time and UDP bulk best effort.
 3. The method according to claim 2 comprising a further step of:
 - 25 delivering the traffic downstream according to the quality of service specified by said one of a plurality of classes.
 4. The method according to claim 1 comprising a further step of:
 - determining a port number which the traffic uses, and
 - 30 classifying the traffic into one of a plurality of classes which includes two default classes, TCP interactive and UDP low latency, according to the port number and preset criteria of traffic characteristics which includes TCP protocol type, UDP protocol type, the length of packets in the traffic and the number of consecutive packets in the traffic, said plurality of classes specifying respective levels of the quality of service with which to condition the traffic.

5. The method according to claim 1 wherein the quality of service is expressed by a plurality of parameters which includes a priority level of delivery of the traffic downstream and the bandwidth assigned to the traffic.

5

6. The method according to claim 1 comprising a further step of: informing one or more downstream nodes the results of monitoring and classifying.

10 7. The method according to claim 6 wherein the step of informing comprises a step of:

propagating the relevant information in a packet which is inserted into the traffic.

15 8. In a packet data network for multimedia traffic having one or more nodes, a method of dynamically conditioning traffic comprising steps of:

continuously monitoring the traffic as to its traffic characteristics at a source host from which the traffic is generated;

20 classifying, at the source host, the traffic into one of a plurality of classes which includes two default classes, TCP interactive and UDP low latency, according to preset criteria of traffic characteristics which includes TCP protocol type, UDP protocol type, the length of packets in the traffic and the number of consecutive packets in the traffic, said plurality of classes specifying respective levels of the quality of service with which to condition the traffic; and

25 25 changing the TCP interactive class to one or the other of TCP bulk guaranteed and TCP bulk best effort, if the traffic contains a series of long packets without an intervening series of two or more short packets.

30 9. The method according to claim 8 comprising a further step of: changing the UDP low latency class to one or the other of UDP real time and UDP bulk best effort.

10. The method according to claim 9 comprising a further step of: delivering the traffic downstream according to the quality of service specified by said one of a plurality of classes.
- 5 11. The method according to claim 8 comprising a further step of: informing one or more downstream nodes the results of monitoring and classifying.
- 10 12. The method according to claim 11 wherein the step of informing comprises a step of:
propagating the relevant information in a packet which is inserted into the traffic.
- 15 13. The method according to claim 8 comprising a further step of:
determining a port number which the traffic uses, and
classifying the traffic into one of a plurality of classes which includes two default classes, TCP interactive and UDP low latency, according to the port number and preset criteria of traffic characteristics which includes TCP protocol type, UDP protocol type, the length of
20 packets in the traffic and the number of consecutive packets in the traffic, said plurality of classes specifying respective levels of the quality of service with which to condition the traffic.




Fig 1





Fig 2



Router or Server Integrated
Traffic Conditioning




Fig 3




Fig 5

Discard all packets except for propagating class info.




Fig 6

5/6




Fig 7

6/6




Fig 8

INTERNATIONAL SEARCH REPORT

International Application No.

PCT/CA 97/00936

A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 H04Q11/04 H04L12/56

According to International Patent Classification(IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	NEWMAN P ET AL: "FLOW LABELLED IP : A CONNECTIONLESS APPROACH TO ATM" PROCEEDINGS IEEE INFOCOM '96, vol. 3, 24 March 1996, pages 1251-1260, XP000622261 * sections 5.2, 5.6, 6 * * Table 1 *	1-13
A	DAMASKOS S ET AL: "THE BERKOM-II FLOWSPEC AND ITS MAPPING ON ATM" SUPERCOMM/ICC '94., 1 May 1994, pages 141-146, XP000438899 see page 141, column 1, line 34 - page 141, column 2, line 4 see page 141, column 2, line 12 - line 31 * section 3.1.3.1 * see page 144, column 2, line 1 - line 27 ---	1,8
		-/-

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance, the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "Z" document member of the same patent family

1

Date of the actual completion of the international search	Date of mailing of the international search report
13 March 1998	25/03/1998
Name and mailing address of the ISA European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl. Fax: (+31-70) 340-3016	Authorized officer Gregori, S

INTERNATIONAL SEARCH REPORT

International Application No

PCT/CA 97/00936

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	EP 0 706 297 A (IBM) 10 April 1996 see page 2, line 53 - line 56 see page 3, line 35 - line 50 -----	1,8

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

In tional Application No

PCT/CA 97/00936

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0706297 A	10-04-96	NONE	